

(12) UK Patent Application (19) GB (11) 2 379 764 (13) A

(43) Date of A Publication 19.03.2003

(21) Application No 0214263.6

(22) Date of Filing 20.06.2002

(30) Priority Data

(31) 09896019 (32) 29.06.2001 (33) US

(71) Applicant(s)

Hewlett-Packard Company
(Incorporated in USA - Delaware)
3000 Hanover Street, Palo Alto,
California 94304, United States of America

(72) Inventor(s)

Tse-Huong Choo
Scott Alan Leerssen
Joubert Berger

(74) Agent and/or Address for Service

Carpmaels & Ransford
43 Bloomsbury Square, LONDON,
WC1A 2RA, United Kingdom

(51) INT CL⁷

G06F 1/00 12/14

(52) UK CL (Edition V)

G4A AAP A23E

(56) Documents Cited

EP 0768594 A1 WO 2002/061554 A

WO 2002/061553 A WO 2002/061552 A

WO 2002/050644 A2

Lin A & Brown R, "The application of security policy to
role-based access control...", Computer
Communications, v23, n17, pp 1584-1593, Nov 2000,
ISSN 0140-3664

(58) Field of Search

UK CL (Edition T) G4A

INT CL⁷ G06F

Other: Online: JAPIO, EPODOC, WPI, TDB, INSPEC,
XPESP

(54) Abstract Title

File system mandatory access control

(57) In one embodiment, the present invention is a computer system including compartments implemented on an operating system. A database 203 contains access rules defining which compartments are authorized to access particular file resources. A kernel module receives a system call to access a file from a user space application 201 belonging to a compartment. A security module 202 determines whether the user space application is authorized to access the file utilizing access rules stored in the database.

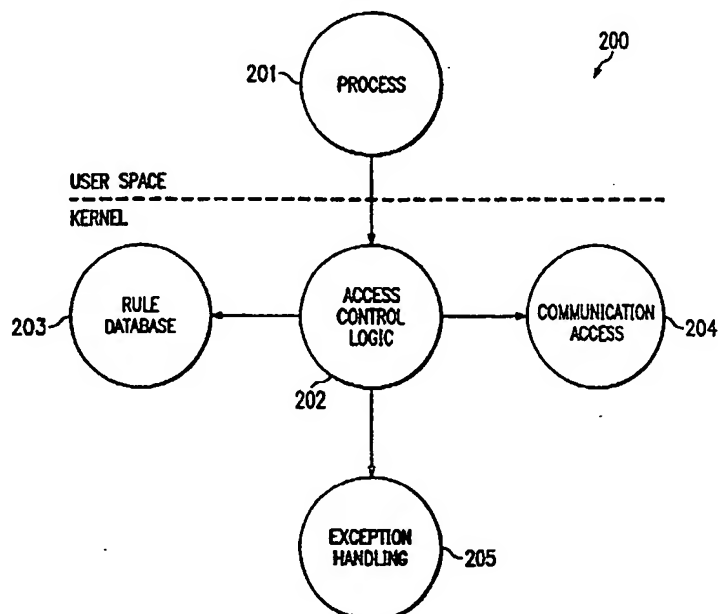


FIG. 2

GB 2 379 764 A

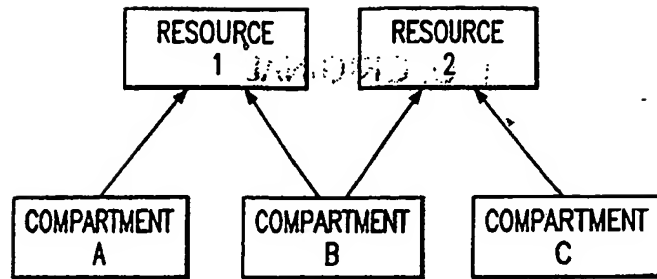


FIG. 1

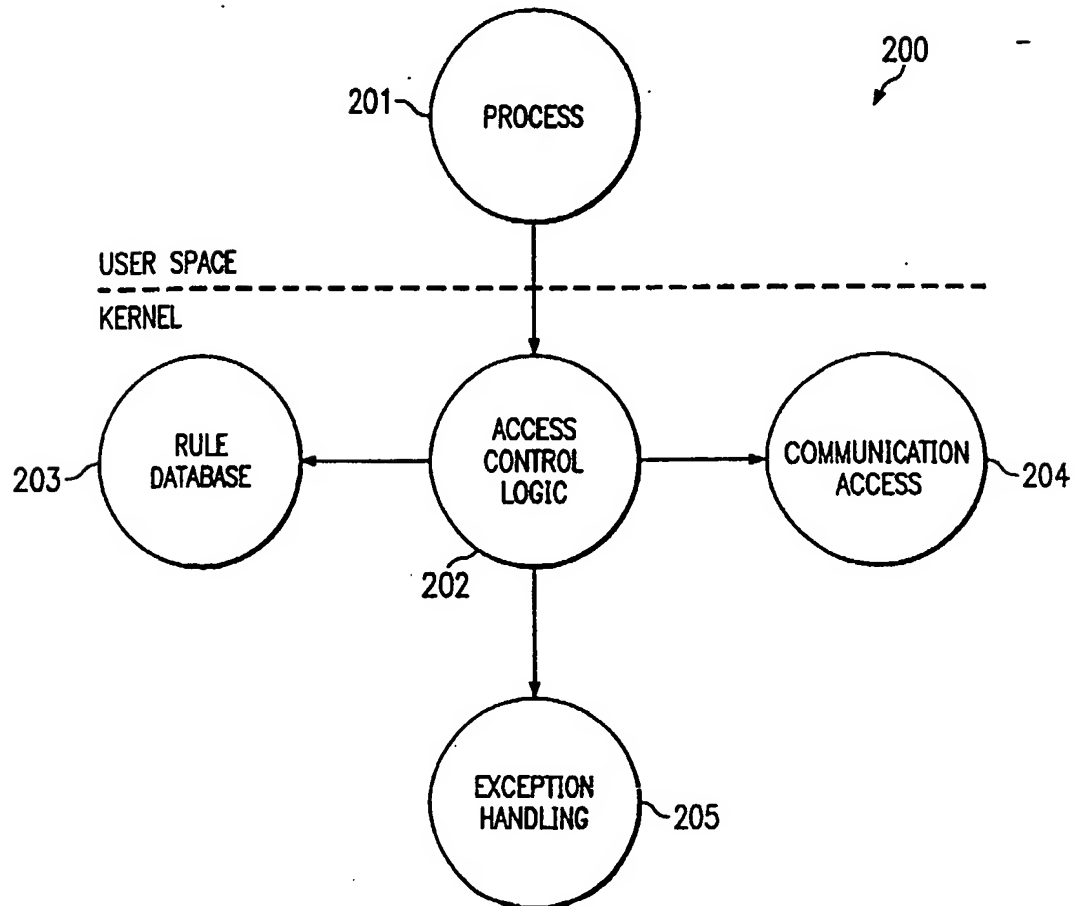
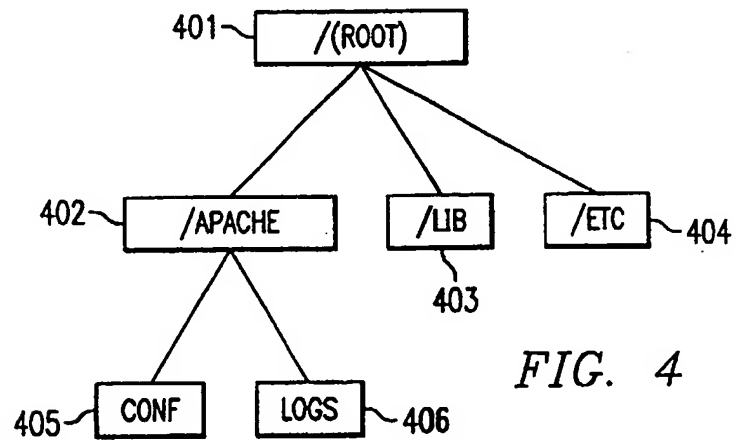



FIG. 2



SYSTEM AND METHOD FOR FILE SYSTEM MANDATORY ACCESS CONTROL**RELATED APPLICATIONS**

This application is related to concurrently filed and commonly assigned U.S. Patent Application Serial No. _____, entitled, "SYSTEM AND METHOD FOR MANAGEMENT OF COMPARTMENTS IN A TRUSTED OPERATING SYSTEM," which is hereby incorporated herein by reference.

TECHNICAL FIELD

The present invention is directed to a system and method for computer containment and more particularly to a system and method for restricting access to files by processes.

BACKGROUND

Containment refers to restrictions on a computer system which prevent user-space applications from performing certain actions. In particular, containment is achieved by forcing a large untrusted application to utilize a smaller trusted application to perform certain actions. By forcing the larger application to do so, the smaller application may ensure that the larger application does not perform undesirable actions, such as interfering with other applications.

One aspect of containment is restricting access to files. For example, it may be advantageous to restrict access to a configuration file, since the configuration file may be utilized to breach the security of the system. Likewise, it is advantageous to prevent most processes from being able to read or write to files containing password information.

To restrict access to files, known trusted operating systems associate access information with each file stored on a file system. Specifically, the file structure is modified to include an additional permission data structure with each file. The permission data structure contains essentially a list of identifiers with each identifier specifying a group of processes that are allowed to access the respective file. When a process attempts to access a particular file, the process performs a system call to the kernel. The identifier of the process is obtained by the kernel routine associated with the system call. The kernel routine accesses the file by reading the list of identifiers. A logical comparison is made between the identifier received from the process and the list of identifiers. If a match is found, the kernel routine performs the access operation (e.g., opening the file). If no match is found, the kernel routine does not perform the access operations and, instead, returns an exception (e.g., error message).

Although associating such a data structure with each file does restrict certain processes from accessing certain files, this approach is problematic in many respects. First, the amount of permission data is large, because file systems of ordinary complexity typically contain thousands of files. Secondly, the task of synchronizing permission data with file creation and file deletion is challenging. For example, many processes may create and delete files during their operation. If permission data is created or modified for each file operation, system performance is significantly degraded. Moreover, if permission data is also maintained by a

system administrator, system administration is quite cumbersome when the number of files exceeds a small number.

5 It shall be appreciated that associating the additional data structure with each file causes the file system format to be incompatible with other file system formats. In particular, this approach is incompatible with the file system formats utilized by traditional UNIX operating systems. Thus, once data is stored in the above format, well-known applications and utilities cannot be utilized with the preceding access limiting file structure.

SUMMARY OF THE INVENTION

In one embodiment, the present invention is related to a computer system including compartments implemented on an operating system. A database contains access rules with said access rules defining which compartments are authorized to access particular file resources. A kernel module receives a system call to access a file from a user space application belonging to a compartment. A security module determines whether said user space application is authorized to access said file utilizing access rules stored in said database.

BRIEF DESCRIPTION OF THE DRAWING

FIGURE 1 depicts a block diagram example of compartments according to the prior art.

FIGURE 2 depicts an exemplary system that utilizes compartments to provide containment according to embodiments of the present invention.

5 FIGURE 3 depicts another exemplary system that utilizes compartments to provide containment according to embodiments of the present invention.

FIGURE 4 depicts an exemplary file system to which access is restricted according to embodiments of the present invention.

DETAILED DESCRIPTION

Compartments refer to groups of processes or threads which are limited to accessing certain subsets of system resources of a computer system. FIGURE 1 depicts a block diagram example of compartments. System 100 includes two subsets of system resources (resource 1 and resource 2). System 100 also includes three compartments (designated compartments A, B, and C). Compartment A is only permitted to access the system resources associated with resource 1. Compartment C is only permitted to access the system resources associated with resource 2. Compartment B is permitted to access the system resources associated with both resource 1 and resource 2. As an example, if a process is designated as belonging to compartment A, the process would be allowed to access resource 1 but would be prevented from accessing resource 2.

According to embodiments of the present invention, by utilizing compartments, the security of a computer system may be enhanced through mandatory access control. Mandatory access control refers to access control that a process cannot override. By utilizing mandatory access control, a breach of security in one compartment will not effect resources associated with another compartment. Specifically, if the security of an application operating in compartment A is compromised, the breach of security is limited to a subset of system resources. For example, resource 1 may include system resources associated with receiving TCP/IP packets without including system resources used to send TCP/IP packets. Instead, the system resources used to send TCP/IP packets may be assigned to resource 2. If an application in compartment A is compromised by a buffer-overflow attack, the application could not be utilized to launch a denial of service attack against another web-resource. The application could not launch such an attack, since it is not permitted to access system resources associated with sending TCP/IP packets.

In embodiments of the present invention, any number of system resources may be organized according to compartment access control. For example, system resources associated with TCP/IP networking, routing tables, routing caches, shared memory, message

queues, semaphores, process/thread handling, and user-id (UID) handling may be limited by utilizing compartments according to embodiments of the present invention.

FIGURE 2 depicts exemplary system 200 that illustrates how compartments may be implemented according to embodiments of the present invention. System 200 includes process 201 that is associated with a compartment. Process 201 executes code in user-space, i.e. a hardware-enforced operating mode that limits the operations of process 201. Process 201 may include code that is operable to attempt to access a protected resource (e.g., opening a certain file) according to a compartment scheme. Process 201 performs a system call to the kernel of the operating system. The system call includes transferring control to access control logic 202. Access control logic 202 receives a compartment identifier or tag of process 201. Access control logic 202 utilizes the compartment identifier to search rule database 203 to determine whether the compartment associated with process 201 is permitted access to the particular resource. If access is permitted by the rules contained in rule database 203, access control logic 202 transfers processing control to communication access module 204 that performs the software operations to access the resource. If access is not permitted, access control logic 202 transfers processing control to exception handling module 205. Exception handling module 205 may return an exception (e.g., an error message) to process 201 and/or it may stop the operations of process 201.

System 300 of FIGURE 3 depicts another exemplary system that utilizes compartments to provide containment. System 300 includes a plurality of compartments. In this example, WEB compartment 301, FTP compartment 302, and SYSTEM compartment 303 are shown. Each compartment is associated with various executing processes or threads. The processes of the compartments are limited to accessing system resources according to the rules stored in rule database 316. Rule database 316 may include various components or modules for the various types of resources. Rule database 316 may comprise separate tables for TCP/IP networking resource rules and for file system resource rules. Also, the various components may be stored in different locations. For example, TCP/IP resource rules may be stored in random access memory while file system resource rules may be stored on the file system.

SYSTEM compartment 303 may include processes that facilitate command line utilities 304 to modify the compartments or rules associated with the compartments. Command line utilities 304 may include commands to create or delete a particular compartment. Command line utilities 304 may further include commands to create, delete, and/or modify the rules stored in rule database 316 that limit access to system resources.

Command line utilities 304 may further include commands to execute a process in a specific compartment. For example, a command may be utilized to execute an HTTP web server application in WEB compartment 301. The command causes a thread to be created. The command also creates an entry in the thread registry of the kernel (not shown). The thread is associated with a unique identifier. Also, the thread is associated with the identifier of WEB compartment 301. When the particular thread makes systems calls to the kernel to access system resources, the kernel utilizes the unique thread identifier to determine the compartment identifier. The kernel then determines whether the particular thread is authorized to access the requested resource. It shall be appreciated that this approach is quite advantageous, because this approach requires no modification to the application being executed. Thus, the exemplary compartment approach described herein allows the security of ordinary platforms to be upgraded to include access control without requiring appreciable modification of user-space application code.

In the example of FIGURE 4, command line utilities 304 access the kernel via kernel modules 322. Routines of kernel modules 322 advantageously perform the actual manipulation (e.g., addition, modification, or deletion) of the respective objects as desired by the particular commands. Further examples of compartment manipulation via command line utilities are disclosed in U.S. Patent Application Serial No. _____, entitled, "SYSTEM AND METHOD FOR MANAGEMENT OF COMPARTMENTS IN A TRUSTED OPERATING SYSTEM," which has been incorporated herein by reference.

The kernel of system 300 includes a plurality of modules. Certain modules are accessed by the various compartments via system calls. For example, processes operating in either WEB compartment 301 or FTP compartment 302 may communicate with processes

operating on other systems via the Internet by utilizing system calls to routines of TCP/IP networking module 306. Socket communication may occur via UNIX domain sockets module 308. Interprocess communication module 310 includes kernel routines to facilitate communication between processes via shared memory, stacks, semaphores, and/or the like.

5 Interprocess communication module 310 may also facilitate spawning or forking new processes. File access module 312 may facilitate access to files on a file system. For example, file access module 312 may facilitate opening, closing, reading from, writing to, deleting, renaming files, and/or the like. Other kernel modules may be provided via other subsystems module 314.

10 Each of the kernel modules advantageously interacts with security module 320. Security module 320 enforces the compartment scheme to prevent unauthorized access to system resources. Security module 320 utilizes device configuration module 318 and rule database 316 to facilitate compartment limitations. Security module 320 is capable of determining which resources are available to system 300 via device configuration module 318.

15 Security module 320 further receives identification of a compartment and identification of a system resource to be accessed from a routine of a kernel module. Security module 320 searches rule database 316 to locate an applicable rule. Security module 320 permits or disallows access upon the basis of an applicable rule, or upon the basis of a default rule if no applicable rule is located.

20 It shall be appreciated that system 300 is an exemplary system. The present invention is not limited to any particular compartment or containment scheme. Specifically, numerous approaches may be utilized to prevent processes belonging to a compartment from accessing system resources. For example, access control may be implemented at the user-level via several techniques. A `strace()` mechanism may be utilized to trace each system call of a given

25 process. The `strace()` mechanism examines each system call and its arguments. The `strace()` mechanism either allows or disallows the system call according to rules defined in a rule database. As another example, system call wrapping may be utilized. In system call wrapping, wrapper functions, using a dynamically linked shared library, examine system calls and arguments. The wrapper functions also either allow or disallow system calls according to

rules defined in a rule database. User-level authorization servers may be utilized to control access to system resources. User-level authorization servers may control access to system resources by providing a controlled data channel to the kernel.

In embodiments of the present invention, access to files by processes is restricted by rules based on process compartments. Reference is now made to FIGURE 4 that depicts exemplary file system 400 to which access is controlled by rules based on process compartments. File system 400 is organized according to a subdirectory structure. The highest component of file system 400 is the root directory (referred to as root 401). Underneath root 401, several subdirectories are shown including /apache 402, /lib 403, /etc 404. It shall be appreciated that any number of subdirectories could exist at any level of file system 400. However, the number of subdirectories shown in FIGURE 4 is limited to aid the reader's understanding of embodiments of the present invention. Additionally, several subdirectories are shown underneath /apache 402 (/apache/conf 405 and /apache/logs 406). As is well known in the art, the pathname to a file in a subdirectory is given by the various subdirectories. For example, the pathname for the file "/apache/conf/user0146.logs" is /apache/conf 406. The pathname and filename may be passed to a function or a system call to perform various access operations such as opening the file, reading from the file, writing to the file, renaming the file, deleting the file, and/or the like.

TABLE I, below, sets forth a number of exemplary rules that may be included in database 316 to control access to file system 400 consistent with the teachings of the present invention:

TABLE I

Rule No.	Compartment	Pathname	Access
1	WEB	/apache/conf	READ
2	WEB	/apache/logs	READ, WRITE
3	WEB	/	NONE (no access)

Rule No.	Compartment	Pathname	Access
4	SYSTEM	/	READ, WRITE

The rules of TABLE I define the permissions given to any process belonging to WEB compartment 301 and SYSTEM compartment 303 to access files within root directory 401 and files within the /apache/conf 405 and /apache/logs 406 subdirectories. For example, a process that belongs to WEB compartment 301 is permitted to read any file within /apache/conf 405 and is allowed to read or write to any file within /apache/logs 406. However, processes belonging to WEB compartment 301 are not permitted any access to files within root directory 401. A process in SYSTEM compartment 303 is permitted read and write access to files within root directory 401.

The rules set forth in TABLE I may be stored in database 316 in any form. However, it is advantageous to store the rules in a manner that parallels the subdirectory structure of file system 400. For example, database 316 may include a series of data structures for each subdirectory of file system 300. The data structures for each subdirectory may contain the rules pertaining to the respective subdirectories. Also, the data structures may form a linked list structure. Specifically, the data structures may contain a pointer to its parent subdirectory and a pointer to each child subdirectory. By organizing the rules in this preferable manner, security module 320 may search the database in an efficient manner by traversing the data structures according to the pathname of the file to be accessed. It shall be appreciated that other mechanisms may be utilized in lieu of a pointer approach. For example, a relational database structure may be utilized to organize rules according to the structure of file system 400.

Additionally, it is advantageous to minimize the number of rules stored in database 316. According to embodiments of the present invention, a default rule may be placed in root directory 401 for compartments. The default rule is applied until another rule is specified at a data structure associated with lower subdirectory. The specific rule in the data structure associated with the lower subdirectory is applied to every child subdirectory thereafter until another rule is located. According to the exemplary rules given in TABLE I, the default rule

for a process belonging to WEB compartment 301 is no access. More specific rules are provided for /apache/conf 405 and /apache/logs 406. By applying this approach, a process belonging to WEB compartment 301 is allowed access to read from every file in /apache/conf 405 and every child subdirectory associated with /apache/conf 405. Likewise, a process
5 belonging to WEB compartment 301 is allowed access to read from and write to every file in /apache/logs 406 and every child subdirectory associated with /apache/logs 406.

According to embodiments of the present invention, security module 320 determines which rules apply based on the compartment identifier of the process. If no rules are located in rule database 316, access is permitted by default. If one or more rules apply, security
10 module 320 preferably utilizes the most specific rule. Specifically, security module 320 first examines the rules to determine whether a specific rule applies to the particular file. If such a rule is located, it is applied. If not, security module 320 examines the lowest subdirectory associated with the file that is defined by the pathname. If a rule is provided for that subdirectory, it is applied. If not, security module 320 successively searches for a rule at each
15 higher parent subdirectory until a rule is located or root directory 401 is reached.

For example, a process belonging to WEB compartment 301 may attempt to read /apache/conf/httpd.conf. A number of rules (Rules 1, 2, and 3) exist for WEB compartment 301. Accordingly, the most specific rule is applied. The rule pertaining to the lowest subdirectory, /apache/conf 405, is applied, i.e. Rule 1, because no rule explicitly exists for
20 apache/conf/httpd.conf. Security module 320 permits access on the basis of Rule 1. Later, the same process belonging to WEB compartment 301 may attempt to write to /apache/conf/httpd.conf. As discussed, Rule 1 applies. In this case, security module 320 does not permit access to the file, because only READ access is permitted by Rule 1.

The same process belonging to WEB compartment 301 may attempt to write to
25 /etc/passwd. A number of rules (Rules 1, 2, and 3) exists for WEB compartment 301. A specific rule is not provided for the file. Accordingly, security module 320 examines the lowest subdirectory defined by the pathname. No rule applies for /etc 404 for WEB compartment 301. Security module 320 searches the parent of /etc 404 which is root

directory 401. Security module 320 locates Rule 3 (no access) which is associated with root directory 401. Accordingly, access is not permitted.

It shall be appreciated that embodiments of the present invention provide several advantages. First, the use of a database to retain access information related to compartments
5 greatly simplifies security management. Specifically, it is not necessary to apply and validate access information to each file. Synchronization issues are significantly reduced, since access information need not be modified for each additional or deleted file. The amount of access information is significantly reduced, because rules are based on subdirectories instead of based on individual files. Structuring the database of rules to parallel the subdirectory structure of
10 the file system allows for efficient access to rules of the database by the kernel. Also, structuring the database in this manner simplifies maintenance of rules by a system administrator. Additionally, it shall be appreciated that embodiments of the present invention are compatible with known file system formats. Specifically, embodiments of the present invention may be implemented without modifying the file structure of files, because a database
15 is utilized that is distinct from the files. Accordingly, embodiments of the present invention allow platforms to implement security procedures without requiring modification of the user-space applications or modification of their file systems.

WHAT IS CLAIMED IS:

1. A computer system for controlling access to certain files by processes, said computer system comprising:
 - compartments implemented on an operating system;
 - 5 a database containing access rules, said access rules defining which compartments are authorized to access particular file resources;
 - a kernel module for receiving a system call to access a file from a user space application belonging to a compartment; and
 - a security module for determining whether said user space application is authorized to
 - 10 access said file utilizing access rules stored in said database.
2. The computer system of claim 1 wherein said database is stored on a common file system with said particular file resources.
3. The computer system of claim 1 wherein each compartment is assigned a unique identifier.
4. The computer system of claim 1 wherein at least one access rule in said database defines whether any process belonging to a particular compartment is permitted to access a plurality of files within at least one subdirectory.
5. The computer system of claim 1 wherein at least one access rule in said database defines whether any process belonging to a particular compartment is permitted to access a particular file.

6. The computer system of claim 1 wherein said security module comprises at least one kernel routine.

7. The computer system of claim 1 wherein said particular file resources are maintained on a file system possessing a subdirectory-based structure and wherein said database organizes access rules according to said subdirectory-based structure.

8. The computer system of claim 7 wherein a default access rule is stored at a database location associated with a root directory of said subdirectory-based structure.

9. The computer system of claim 7 wherein specific access rules are stored at database locations associated with subdirectories of said subdirectory-based structure.

10. The computer system of claim 9 wherein said security module is operable to receive a path identifier of the file from said user space application with said path identifier including a plurality of subdirectory identifiers.

11. The computer system of claim 10 wherein said security module is operable to apply an access rule according to a lowest subdirectory identifier of said plurality of subdirectory identifiers.

12. The computer system of claim 10 wherein said security module is operable to respectively search said database according to each higher subdirectory identifier of said plurality of subdirectory identifiers, when an access rule is not located according to a lower subdirectory identifier.

13. The computer system of claim 1 wherein said security module is operable to permit access when no access rule is located.

14. A method for controlling access to a file by a process, said method comprising:
receiving a request from said process to access said file, said process being associated
with a compartment implemented on an operating system;
determining an identifier of said compartment; and
5 searching for access rules on a database, said database containing access rules defining
whether processes associated with particular compartments are permitted to access certain file
resources.
15. The method of claim 14 wherein said file is stored on a file system that
possesses a subdirectory structure, and wherein said database is structured to retain access
rules in a hierarchical manner that parallels the subdirectory structure of said file system.
16. The method of claim 15 wherein said access rules includes at least one access
rule that allows a process associated with said compartment to access a plurality of files
associated with a particular subdirectory.
17. The method of claim 15 wherein a default access rule stored in said database is
associated with a root directory of said file system.
18. The method of claim 17 wherein specific access rules are stored in said
database and said specific access rules are associated with subdirectories of said file system.
19. The method of claim 14 wherein said request includes a filename containing a
path identifier, said path identifier specifying a plurality of subdirectories, wherein said step of
searching includes the sub-steps of:

- 5 (a) searching said database according to a lowest subdirectory of said plurality of subdirectories for an access rule applicable to said compartment;
- (b) when an access rule is found in step (a), proceeding to step (e);
- (c) searching said database according a next higher subdirectory of said plurality of subdirectories for an access rule applicable to said compartment;
- 10 (d) repeating step (c) until the first event of the following events occurs:
- (i) an access rule applicable to said compartment is located;
- (ii) said database is searched according to a root directory.
- (e) when an access rule applicable to said compartment is located, providing access to said file when said access rule applicable to said compartment allows access.

20. The method of claim 19 wherein said step of searching further comprises:

- (f) when an access rule applicable to said compartment is not located, providing access to said file.

21. The method of claim 14 wherein said database is stored on a same file system as said file.

22. The method of claim 14 wherein said step of searching is performed by a kernel routine of an operating system.

23. The method of claim 14 wherein said database comprises at least one rule that defines whether a process associated with a particular compartment is permitted to access a plurality of files in a particular subdirectory.

24. A computer readable medium including instructions executable by a processor, said computer readable medium comprising:

code for receiving a request from a process associated with a particular compartment to access a particular file, said compartment being associated with an operating system; and

5 code for searching a database containing access rules which define which compartments possess authorization to access certain file resources.

25. The computer readable medium of claim 24 wherein said database comprises at least one rule which defines whether a process associated with a compartment is permitted to access a plurality of files of a subdirectory.

26. The computer readable medium of claim 24 wherein said particular file is stored on a file system that possesses a subdirectory structure, and wherein said database possesses a structure that parallels said subdirectory structure.

27. The computer readable medium of claim 26 wherein said code for receiving receives a filename that possesses a plurality of subdirectory identifiers, and wherein said code for searching searches said database according to a lowest subdirectory of said plurality of subdirectories for an access rule applicable to said compartment.

28. The computer readable medium of claim 26 wherein said code for searching searches said database according to a higher level subdirectory when an access rule applicable to said compartment is not located according to a lower level subdirectory.

29. The computer readable medium of claim 24 further comprising:

code for determining whether said process may access said particular file.

30. The computer readable medium of claim 24 further comprising:
code for denying access to said particular file by said process.



INVESTOR IN PEOPLE

Application No: GB 0214263.6
Claims searched: 1-30

Examiner: Geoff Western
Date of search: 13 January 2003

Patents Act 1977 : Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X,E	1,14,24, at least	WO 2002/061554 A (NORMAN et al) N.b. pp 4-8, p15 lines 21-26, p16
X,E	1,14,24, at least	WO 2002/061553 A (CHOO et al) N.b. pp 4-8, p15 lines 23-29, p16
X,E	1,14,24, at least	WO 2002/061552 A (NORMAN et al) N.b. pp 4-8, p15 lines 23-29, p16
X,E	1,14,24, at least	WO 2002/050644 A2 (SUN) N.b. pp 12-13
X	1,14,24, at least	EP 0768594 A1 (DATA GENERAL) N.b. pp 1, 6, 7
X	1,14,24, at least	Lin A & Brown R, "The application of security policy to role-based access control and the common data security architecture", Computer Communications, v23, n17, pp 1584-1593, Nov 2000, ISSN 0140-3664

Categories:

X Document indicating lack of novelty or inventive step	A Document indicating technological background and/or state of the art.
Y Document indicating lack of inventive step if combined with one or more other documents of same category.	P Document published on or after the declared priority date but before the filing date of this invention.
& Member of the same patent family	E Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^v:

G4A

Worldwide search of patent documents classified in the following areas of the IPC⁷:

G06F

The following online and other databases have been used in the preparation of this search report:

Online: JAPIO, EPODOC, WPI, TDB, INSPEC, XPESP